

3-2005

An Historical Analysis of Factors Contributing to the Emergence of the Intrusion Detection Discipline and its Role in Information Assurance

James L.M. Hart

Follow this and additional works at: <https://scholar.afit.edu/etd>



Part of the [Risk Analysis Commons](#)

Recommended Citation

Hart, James L.M., "An Historical Analysis of Factors Contributing to the Emergence of the Intrusion Detection Discipline and its Role in Information Assurance" (2005). *Theses and Dissertations*. 3814.
<https://scholar.afit.edu/etd/3814>

This Thesis is brought to you for free and open access by the Student Graduate Works at AFIT Scholar. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of AFIT Scholar. For more information, please contact AFIT.ENWL.Repository@us.af.mil.



**AN HISTORICAL ANALYSIS OF FACTORS CONTRIBUTING TO THE
EMERGENCE OF THE INTRUSION DETECTION DISCIPLINE AND ITS
ROLE IN INFORMATION ASSURANCE**

THESIS

James L. M. Hart, Captain, USAF

AFIT/GIR/ENV/05M-06

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

AFIT/GIR/ENV/05M-06

**AN HISTORICAL ANALYSIS OF FACTORS CONTRIBUTING TO THE
EMERGENCE OF THE INTRUSION DETECTION DISCIPLINE AND ITS
ROLE IN INFORMATION ASSURANCE**

THESIS

Presented to the Faculty

Department of Systems and Engineering Management

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the
Degree of Master of Science in Information Resource Management

James L. M. Hart, BS

Captain, USAF

March 2005

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

**AN HISTORICAL ANALYSIS OF FACTORS CONTRIBUTING TO THE
EMERGENCE OF THE INTRUSION DETECTION DISCIPLINE AND ITS
ROLE IN INFORMATION ASSURANCE**

James L. M. Hart, BS

Captain, USAF

Approved:

 /signed/
David D. Bouvin, Capt, USAF, PhD (Chairman)

15 March 2005
Date

 /signed/
Kevin L. Elder, PhD (Member)

15 March 2005
Date

 /signed/
Dennis D. Strouble, PhD (Member)

15 March 2005
Date

Abstract

In 2003, Gartner, Inc., predicted the inevitable demise of the intrusion detection (ID) market, a major player in the computer security technology industry. In light of this prediction, IT executives need to know if intrusion detection technologies serve a strategic purpose within the framework of information assurance (IA). This research investigated the historical background and circumstances that led to the birth of the intrusion detection field and explored the evolution of the discipline through current research in order to identify appropriate roles for IDS technology within an information assurance framework. The research identified three factors contributing to the birth of ID including increased procurement and employment of resource-sharing computer systems in the DoD, a growing need to operate in an open computing environment while maintaining security and the unmanageable volume of audit data produced as a result of security requirements. The research also uncovered six trends that could be used to describe the evolution of the ID discipline. Finally, the research outlined three roles suitable for IDS to fulfill within the IA framework.

Acknowledgments

I would first like to thank my beautiful wife and children for their unending support during this thesis effort. I can't wait to focus my off-time on being a husband and father again. Additionally, I'd like to thank my faculty advisor for his guidance and patience as I muddled through the graduate education process. Finally, I'd like to offer my thanks and congratulations to the rest of my colleagues and faculty in the IRM program. I feel very fortunate being able to share a learning environment with men and women of your caliber.

James L. M. Hart

Table of Contents

	Page
Abstract.....	iv
Acknowledgments.....	v
Table of Contents.....	vi
List of Figures	ix
I. Introduction	1
Background.....	1
Problem Statement.....	5
Research Objectives	5
Research Focus	6
Investigative Questions	6
Methodology.....	6
Limitations.....	7
Implications	7
Summary.....	8
II. Review of Relevant Literature	9
Chapter Overview.....	9
The Emergence of the Problem	9
Developing a Plan of Attack.....	18
Audit in the 1970's	21
Intrusion Detection Is Born	22
An Intrusion Detection Model.....	26
IDS in the late 1980's	28

	Page
Surveys of IDS Research.....	35
IDS Market Trends	40
Information Assurance	42
Summary.....	46
III. Methodology	47
Chapter Overview.....	47
The Importance of History	47
The Historical Method.....	48
Justifying an Historical Analysis.....	52
Limitations of the Historical Method	53
Summary.....	54
IV. Findings	55
Chapter Overview.....	55
Addressing the Investigative Questions	55
<i>Investigative Question #1: What factors led to the emergence of intrusion detection as a discipline?</i>	56
<i>Investigative Question #2: How has IDS research evolved since its inception?</i>	57
<i>Investigative Question #3: What specific roles, if any, are IDS's suited to fulfill within a holistic information assurance program?</i>	58
A Synthesis.....	59
Summary.....	60
V. Summary, Implications, and Recommendations.....	62
Summary.....	62
Practical Implications	63

	Page
Recommendations for Future Research.....	64
Bibliography	66
Vita	70

List of Figures

	Page
Figure 1: CSI/FBI Computer Crime Survey Respondents Employing IDS	3
Figure 2: General Cases of Threats.....	25
Figure 3: Classical Detection Theory	39
Figure 4: Worldwide IDS/IPS Product Revenue Breakdown by Year	41
Figure 5: Information Assurance	43
Figure 6: Defense in Depth Strategy Focus Areas.....	45
Figure 7: The Cascade	49
Figure 8: Mason-O'Brien Synthesis	52

AN HISTORICAL ANALYSIS OF FACTORS CONTRIBUTING TO THE EMERGENCE OF THE INTRUSION DETECTION DISCIPLINE AND ITS ROLE IN INFORMATION ASSURANCE

I. Introduction

Background

The concept of security in today's world has become pervasive in all facets of life. A large portion of the American workforce interfaces daily with computer systems through the use of email applications, word processing applications, spreadsheets, and presentation software. With such widespread use of computing equipment, the idea of computer security has also become embedded in our society. Ever since the Michelangelo virus of 1992, when computer users emptied the store shelves of commercial antiviral software (Anonymous, 2004), the developed world has cultivated a keen awareness of computer security issues.

While most of us comprehend the concept of computer security, we still must wonder what the components are that make up a secure computer system? Not too long ago, it seemed that the average PC owner needed only an updated virus scanner to keep their systems safe from malicious computer code. Now, with the widespread use of "always-on" network connections, the average home computer is just as susceptible to attack by hackers as big corporations used to be. Furthermore, if left unprotected, these

systems can be, in effect, taken over and employed as platforms for mounting attacks against other systems. In light of this development, computer security experts are advising the average user to invest in a more robust security strategy that entails traffic filtering as well as virus protection. If this is the case for the average computer user, what does it bode for corporate users who integrate their business operations using computing technologies. As processing speeds and computational power increase, and as users progressively realize greater levels of interconnection, cyber-criminals are able to practice their trade with more ease and anonymity than ever before. Furthermore, as organizations continue to capitalize on the efficiencies offered by the democratization of technology, it seems that the efficiencies are often overshadowed by catastrophic losses due to computer crime. This leaves our information technology managers to ponder which tools at their disposal are genuinely effective against the threat of computer crime.

The first line of defense against unwanted internet traffic is typically implemented in the form of firewall technologies which filter traffic according to specific rules based on the security policy of a given organization. One author likens a firewall to a security guard in an office building or other secure facility (Rogers, 2002). The guard uses various techniques to screen visitors and decide whether or not they are allowed to enter the facility. The firewall is an excellent tool for filtering unwanted external traffic within a network environment. However, according to the 1999 CSI/FBI Computer Crime Survey, 71 percent of respondents indicated unauthorized accesses by insiders. Couple this information with the ability of skilled malefactors to circumvent firewall technology by exploiting vulnerabilities in commonly used services and you have a significant gap in

security strategy. This is where intrusion detection (ID) potentially provides value.

According to two security experts, “Intrusion detection systems remain the only proactive means of detecting and responding to threats that stem from both inside and outside a corporate network” (Innella & McMillan, 2001).

While intrusion detection systems (IDSs) have been increasingly incorporated into network security strategies since 1999 (Figure 1), the technology has come under fire recently for failing to deliver on its promises of tighter security.

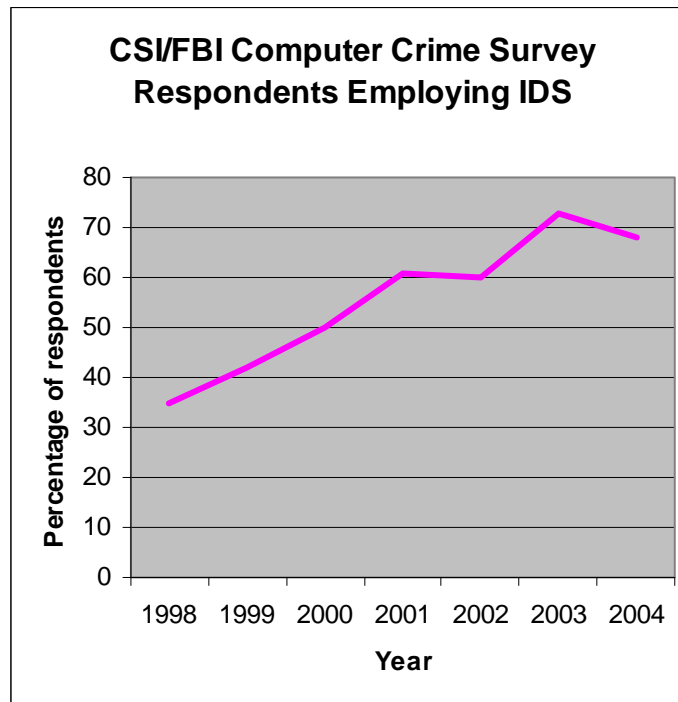


Chart adapted from 2003/2004 CSI/FBI Computer Crime Surveys

Figure 1

In 2003, market research and consulting firm Gartner, Inc. released a scathing criticism of the technology claiming, “IDSs have failed to provide value relative to its costs and will be obsolete by 2005” (Gartner, 2003). While this analysis spurred a backlash by many IDS proponents, the Gartner report proposed a compelling argument for rethinking investment strategies in information security technology. The supporting evidence provided by Gartner included:

- Unmanageable volume of false positives and negatives
- Increased burdens on IS organizations due to full-time monitoring requirements
- A taxing incident-response process
- Inability to monitor traffic rates greater than 600 Mbits/sec

Furthermore, Gartner recommended diverting funds earmarked for IDS to state-of-the-art firewalls that bundle network and application-based defenses in a single solution as well as incorporate deep packet inspection and antivirus capabilities, effectively blocking malicious traffic before it has the opportunity to cause damage.

This criticism came as quite a shock to many in the IT community as the IDS market had earned a significant share of the IT security market starting around 1998. Many large corporations, concerned with respect to protecting their information resources, invested significant capital in implementing commercial IDS products as part of their corporate security strategy. Gartner’s conclusions, while not wholly accepted throughout the IT community, may have caused at least a few IT executives to reconsider their IDS investment. This is especially true when the entire community of IDS users

have complaints consistent with those that Gartner brought to light, including IDS researchers.

Intrusion detection, thus, has suffered from mixed views on its effectiveness in defending network resources. This dichotomy of opinions leads to the purpose of this research and the problem that the researcher aims to shed light on.

Problem Statement

Information systems (IS) and information technology (IT) executives need a strategic understanding of the capabilities, limitations, and underlying theory of detection for IDS technologies in order to make appropriate investment decisions regarding their use. Furthermore, this understanding should be framed in the overarching strategy of information assurance (IA), as IA is the prevailing construct for securing information and information resources employed by IS/IT executives today.

Research Objectives

In order to understand the capabilities and limitations of IDS, it is necessary to comprehend the path that the foundational IDS research has taken to build the existing body of knowledge. Whether or not the discipline or its application has been rendered ineffective, it is important to know from whence the discipline emerged and under what circumstances. Furthermore, it would be beneficial to enumerate the evolution of the IDS discipline thereby providing a strategic perspective on how IDS capabilities have grown in breadth and depth over time. Finally, by enlightening the circumstances by which the IDS discipline emerged and observing the evolutionary patterns that guided development

of the discipline, it may be possible to determine those roles, in a holistic IA program, that this technology is suited to fulfill.

Research Focus

Consequently, the primary focus of this research is to determine what value, if any, the ID function serves with respect to implementing a holistic information assurance program. This will be accomplished by addressing three specific investigative questions.

Investigative Questions

1. What factors led to the emergence of intrusion detection as a discipline?
2. How has IDS research evolved since its inception?
3. What specific roles, if any, are IDS's suited to fulfill within a holistic information assurance program?

Methodology

The proposed methodology for this research effort is the historical analysis. As attested to by Leedy and Ormrod, the focus of historical research is not only to accumulate factual information, but to interpret those facts for some valuable contribution to the body of knowledge (2001). Subsequently, the researcher will apply a systematic process of gathering and critiquing sources of data to obtain enlightenment about the causes, effects, or trends of past events that may result in a better understanding of current or future events.

Limitations

Limitations of this study include a lack of exposure to private sector operations and planning processes on the researcher's part. However, the researcher does have experience in maintaining and managing information technology for the United States Air Force as well as background in the strategic employment of IT. Furthermore, the researcher is attempting to interpret a highly technical topic from a strategic perspective. As such, certain conclusions may not fully consider the range of technical capabilities currently available in the IDS marketplace. Finally, due to the broad range of research within the IDS community, the researcher surveyed primarily seminal works in the discipline and well-cited surveys of available literature. Thus, it's possible that certain perspectives or approaches to intrusion detection have not been represented.

Implications

The implications of this study are important primarily to those employed in the security of computers and computer networks. IS/IT executives will gain the most benefit from the conclusions of this study as it aims to provide guidelines for employing intrusion detection systems within the context of an IA strategy. While these implications may also be cross-utilized in the intrusion detection field as possible directions for further research, the objective is primarily to aid in the formulation and implementation of IT policy encompassing programming, architecting, organizing and budgeting for security and, ultimately, attainment of strategic advantage.

Summary

Chapter One provided background information to frame the context of this research and conveyed the problem that the proposed research methodology will be applied to. This chapter also specified the primary objectives and investigative questions to be addressed throughout the course of the research as well as a brief overview of the historical research methodology. Finally, this chapter introduced notional implications as well as limitations of the research presented in this thesis. In Chapter Two, the researcher presents a review of the literature to address the research question.

II. Review of Relevant Literature

Chapter Overview

Chapter One provided the background and description of the problem addressed by this research as well as a brief synopsis of the proposed historical analysis technique. This chapter explores the available literature to address the investigative questions posed in Chapter One. The discipline of intrusion detection emerged from a greater problem of computer security. As computer systems are primarily employed for the storage, processing and transfer of information, the requirement for computer security is paramount to safeguarding this valuable resource. This chapter examines the backdrop of the general computer security problem from which intrusion detection emerged and then examines the evolution of the IDS discipline from its inception to current research in the field. Finally, this chapter explores the IA construct in order to frame the discussion of IDS within an IA strategy.

The Emergence of the Problem

The discipline of intrusion detection was borne out of the broader topic of computer security. Hence, an analysis of those factors that brought about an interest in computer security should intuitively shed light on the subordinate discipline of intrusion detection.

Many experts consider the Defense Science Board's Task Force on Computer Security report entitled *Security Controls for Computer Systems* to be the first

concentrated effort to address the problem of computer security. This report, commonly referred to as the “Ware report” in honor of the Task Force chairman, was originally published in February of 1970, and was republished in 1979 following declassification. The value of this report is evidenced by its inclusion in the list of seminal papers on computer security compiled by the Computer Security Laboratory within the Computer Science Department at the University of California, Davis.

The Ware report initially traces the emergence of computer security as a serious area of concern to events which occurred during the spring and summer of 1967. This report was prepared ultimately for the United States Department of Defense’s (DoD) Defense Science Board in response to increased anxiety by defense contractors and military operators over the lack of security safeguards present in *resource-sharing systems*, which were being acquired vigorously during this period (Ware, 1970).

Resource-sharing systems essentially dole the resources of a computer system (memory, computational power, peripheral devices) among several concurrent users. In the context of the Ware report, the term included time-sharing, multiprogrammed, remote batch, on-line, multi-access, and multiprocessing systems. The primary difference between the various types was determined by the location of the user relative to the job entry platform when the computational task is being accomplished. Time-sharing, on-line, and multi-access systems required the user to interact with the system while the task, or job, is being performed. Multiprogrammed and remote batch systems completed the job autonomously after the user specified the task and specific criteria for the job. In multiprogrammed and remote batch environments, the system scheduled the job in a

queue and the user must retrieve the results at a later time. Multiprocessing systems were a special case where the system possesses two or more processors which shared the system memory. Typically, resource-sharing systems provided access to users via geographically separated consoles, or terminals connected to the central computer by communication lines (Ware, 1970). These systems were precursors to today's networks and client-server architectures.

In June 1967, the Deputy Director of Defense Research and Engineering (Administration, Evaluation, and Management) contacted the Director of the Advanced Research Projects Agency (ARPA) and formally requested a Task Force be formed to study and recommend hardware and software safeguards that would adequately shield classified information from inadvertent disclosure in a multi-access, resource-sharing computing environment. The Task Force on Computer Security was formed in October following discussions with representatives from academic and industrial communities during the summer and fall of 1967. The Task Force consisted of a Steering Group chaired by Mr. Willis H. Ware from the Rand Corporation as well as a Policy Panel and a Technical Panel reporting to the Steering Group. Mr. Ware was an ex-officio member of both subordinate panels (Ware, 1970).

In the report, the Task Force recognized that the basic problem of machine processing of classified information was not new, even in 1970. The problems had been previously encountered in batch-processing mode and the more recent advances of remote job-entry systems where inadequate security mechanisms had been deliberately or inadvertently circumvented to allow access to sensitive information by unauthorized

users. Still, the mechanisms employed to safeguard these systems had, for the most part, been modified versions of the well-known manual processes of protecting sensitive information. Ware provided an example demonstrating that “the basic principle underlying the security of computer systems has traditionally been that of isolation—simply removing the entire system to a physical environment in which penetrability is acceptably minimized” (Ware, 1970:1).

The Task Force emphasized that the advent of resource-sharing systems, similar in many respects to today’s networks and client-server architectures, had compounded the problem and introduced more complexity to the situation. The isolation technique was no longer robust to safeguarding sensitive information given the impossibility of physically securing systems in which some components, such as user access terminals, are geographically separated.

In explaining the nature of the computer security problem, the Task Force recognized the perspective that the vulnerabilities of resource-sharing systems might be viewed as a trade-off for the efficiencies these systems have to offer (Ware, 1970). This particular view is still pervasive today, but the Task Force felt that it obscured two more fundamental issues. First, the problem of computer security is not limited to one or even a few types of computer systems or configurations; it maintains relevance irrespective of the platform. Ware explained, “...we are really dealing not with system configurations, but with security” (1970:1). The technology had served as a catalyst for shifting the focus from securing the hardware to protecting the **information** independent of the platform. Secondly, the Task Force recognized that resource-sharing systems must be

designed from the very start with a fundamental focus on protecting information not only from other users but from the system itself (Ware, 1970).

Herein lay the value of the Ware report. It was the first serious effort at scoping the problem of computer security. The Ware report established a lexicon for the computing community to begin discussing the issue of security (Anderson, 1972b). At one level of abstraction, this report produced valuable observations within the frame of computer security, but many of these same conclusions are directly relevant to intrusion detection. These conclusions were communicated through a series of recommendations and comments from each panel of the Task Force. They encompassed specific characteristics that each panel believed should be included in a secure computer system. Additionally, these recommendations aided in shaping the efforts of at least one of the seminal researchers in intrusion detection.

While the Policy Panel's recommendations were broad in range and included topics such as personnel roles and responsibilities and information structures in a secure system, the panel also explicitly identified certain system capabilities which were required to maintain security in computing environments. Those capabilities that, at least superficially, are relevant to the field of intrusion detection included system transaction accounting and reliability and auto-testing.

The capability for system transaction accounting required that all significant transactions relevant to security occurring between the users and the system should be documented and automatically time stamped. This capability would then allow for the creation of an audit trail that provided information about file accesses, modifications,

creation, destruction, classification, and reclassification. This audit trail would be made available to security personnel who are, in turn, responsible for maintaining the integrity of the system. Ware also suggested that the activities of system operators, administrators, and maintainer's would be subject to the same level of monitoring, presumably to protect against insider threat issues (Ware, 1970).

The reliability and auto-testing capability required that security controls be redundant as well as self-testing in order to reduce the probability of undetected compromise. In addition to the built-in tests, security personnel should also be charged with periodically inspecting the protection mechanisms of the system. When any of these mechanisms detect a security breach, the system should transition to a degraded mode of operation that ceases the flow of information between the system and the user community. Once the threat of further information leakage has been averted, the system's security personnel are notified and should attempt to resolve the breakdown of the protection mechanisms while adhering to well-documented procedures that maximize security as well as service to the users (Ware, 1970). In making this recommendation, the Task Force implied at least a calculated response capability, if not automated responses to security breaches.

The next section of the Ware report dealt with the technical considerations in maintaining computer security. The Technical Panel, in introducing its recommendations, conceded:

Present technology offers no way to absolutely protect information or the computer operating system itself from all security threats posed by the human beings around it. As a consequence, procedural and administrative safeguards

must be applied in resource-sharing computer centers to supplement the protection available in the hardware and software (Ware, 1970:26).

Even at this early stage in the study of computer security, the panel members admitted that technology alone was not capable of adequately securing the information resident in a computer system.

The Technical Panel, of which James P. Anderson was a member and whose contributions to IDS will be discussed extensively later in this research, presented several technical enhancements that they claimed were necessary and sufficient for ensuring security in a closed environment. The engineering-related recommendations pertaining specifically to the field of intrusion detection were in the areas of supervisor protection, system access control, terminal identification, and certification.

The supervisor protection enhancement provided a mechanism which denied a user program the ability to penetrate the Supervisor, also called the operating system kernel, without the Supervisor detecting it (Ware, 1970). The import of this recommendation is the implication that suspicious or anomalous activity within a system, such as normal users gaining access to privileged processes and functions, should raise attention by security mechanisms or personnel.

Regarding system access control, the Technical Panel recommended a specific enhancement related to denial of access that is relevant to the intrusion detection discipline. The panel recommended that the average user should not be able to gather information regarding the security controls or protection mechanisms emplaced when access is denied him, a common problem in many protection mechanisms. The system

should assume inadvertence on the user's part and assist in identifying any mistakes in negotiating protection mechanisms (Ware, 1970). While this general guideline is important in security design, the correlation to ID resides in the panel's expansion on the idea by further proposing that the system should also continue to log unsuccessful attempts to access sensitive files, processes and system functions, presumably to extract information regarding misuse. This is important today as the use of intrusion detection systems for gathering forensic evidence in prosecuting computer crimes is both a popular and volatile issue. On further expansion, the panel proposed that secure systems should have a means of positively identifying any terminals with which it is communicating and that the system should be able to request terminal identification at any time (Ware, 1970). Again, while this is a generally sound recommendation as far as security is concerned, in the context of intrusion detection this information could prove valuable in correlating penetration attempts.

In the area of certification, the Technical Panel, echoing the Policy Panel's proposal, proposed the generation and use of audit trails to ensure proper usage of the system. However, the technical panel took this one step further by describing specific details that should be incorporated into the system audit (Ware, 1970). While enumerating these details is outside the scope of this research, it is interesting to note this early effort at codifying specific audit data that would aid in maintaining computer security.

Another notable comment about audit trail data made by the technical panel was that special data reduction programs, event-correlation programs, and data-summary

programs will be required by security personnel since a large volume of information will be available through the various logs (Ware, 1970). This early observation foreshadowed the need for separate sensors and data correlation systems that dominate current IDS architectures.

The Technical Panel also echoed the Policy Panel's suggestion for security violation and auto-testing. They proposed that any user activity that violated any security control, should result in the activity being immediately terminated and security personnel notified (Ware, 1970). The panel members apparently advocated automated or active response mechanisms when designing secure systems, especially given the military context for which this report was compiled.

Another interesting observation in the Ware report is its treatment of open and closed systems and the peculiarities of each when discussing security. The panel prefaced their recommendations by acknowledging the existence of two distinct environments for the operation of secure systems—closed environments and open environments. A *closed environment* consists of cleared, or trusted, personnel operating physically protected terminals connected to a physically protected central system over physically protected communication lines. An *open environment* is defined in the Ware report as a “mixture of uncleared users working at unprotected consoles connected to the computing central using unprotected communication circuits and cleared users with protected consoles and protected communication lines” (Ware, 1970:vi).

The Technical Panel asserted that the open computing environment adds significant complexity to the problem of security. The security problem with such an

open system is that it must be capable of negotiating penetration from within and without.

The panel documented specific concerns about the risks inherent when sensitive

(classified) information resides within an open system. These concerns included:

- It is virtually impossible to verify that a large software system is completely free of errors and anomalies
- The state of system design of large software systems is such that frequent changes to the system can be expected
- Certification of a system is not a fully developed technique nor are its details thoroughly worked out
- System failure modes are not thoroughly understood, catalogued, or protected against
- Large hardware complexes cannot be guaranteed error-free (Ware, 1970:26)

While the panel suggested that the collective implementation of their recommendations would result in adequate security for a closed system, the panel could not guarantee the effectiveness of their recommendations within an open environment. Even with a cursory consideration of the security requirements of open systems, the Task Force simply could not predetermine the security implications of open systems. This is mostly due to a lack of experience in dealing with the open environment. After all, in 1970, connectivity such as we enjoy today was unheard of.

As stated previously, the recommendations from this study were the foremost to attack the fledgling issue of computer security; however, even in such an early attempt, the roots of the intrusion detection discipline had begun to materialize.

Developing a Plan of Attack

In 1972, James P. Anderson, a former member of the Technical Panel of the Defense Science Board's Task Force on Computer Security, was called upon to develop a plan of study for the U.S. Air Force to address the emergent issues that were identified by

the Ware report. The Ware report, he acknowledged in the executive summary, “was an important milestone but did not have the impact intended and may have had a negative effect due to its specification of necessary, but not sufficient, criteria for evaluating hardware and software suitable for secure operations” (Anderson, 1972a:4).

Furthermore, Anderson explained that the expected result from the Task Force’s efforts was a recommended research and development program that would lead to the resolution of the computer security problem. While the Ware report did not produce a valid research program to deal with computer security as was intended, the report was the first to properly frame the computer security problem which subsequently allowed further expansion on the topic (Anderson, 1972b). Consequently, Anderson was commissioned to pick up where the Task Force left off.

At the time, it seems that a significant portion of the high level decision makers may have been proponents of the “add-on” approach to security. Anderson and his panel of researchers staunchly disagreed. It was the strong opinion of the panel that solutions to this problem were not going to be solved by augmenting existing systems. In fact, a central theme of the study is the idea that security must be designed into the system from the start. He expanded, “the issue of computer security is one of completeness rather than degree, and a complete system will provide all of the controls necessary for a mixture of all security levels on a single system” (Anderson, 1972a:iv).

The logic for this opinion was that in order to provide defense against a malicious user, the security mechanisms must be designed such that not only are the user’s actions under control but they should be capable of also controlling the various components of

the operating system when it is acting as an intermediary between the user and the information residing in the system (Anderson, 1972a). In general, Anderson's approach to solving this problem was to first formulate a sound theoretical model for a secure computing system and then progress toward an appropriate implementation.

In addition to formulating a development plan for a prototype secure multi-level computer system, Anderson also formulated development plans for supporting technologies. Of particular interest, in the scope of this research, is the proposed development of a Security Surveillance System found in the Exploratory Development Plan. The Exploratory Development Plan was a subset of the overall plan comprised of semi-independent topics which indirectly supported the primary objective of the study.

The objectives of the Security Surveillance System were:

- To detect security-related events (i.e. system behavior which constitutes or precipitates security incidents or violations)
- To collect, record, reduce and analyze data regarding event detections in order to invoke an appropriate compensatory procedure (e.g. exception processor, alarm or correction mechanism)
- To generate reports for security personnel review and damage assessment (Anderson, 1972b:51)

Anderson enumerated the particular areas of such a capability that should be focused on which included instrumentation, measurement, compensatory procedures, reporting and integrity. "Instrumentation," Anderson said, "was a two-fold problem" (Anderson, 1972b:51). What should be detected and how should it be detected? These questions are further complicated by the fact that all system events may be relevant, but not necessarily at the same time or in the same combination. Measurement dealt with how the data is collected, recorded, reduced and analyzed to determine security

implications. Compensatory procedures could be considered what we now refer to as *incident response* and deals with how systems and security personnel react when confronted with a breach in the protection mechanisms. Reporting refers to presenting the collected data such that security personnel are able to intuitively recognize security incidents while also providing flexibility for thorough investigation of an incident. Integrity corresponds to those measures implemented to maintain the security of the surveillance system itself as well as measures present to prevent circumventing the security surveillance system (Anderson, 1972b). According to McHugh, these issues are still at the heart of IDS research today (2001).

Audit in the 1970's

From 1972 – 1980, several computer security research initiatives were undertaken in the public sector with extensive funding from the DoD. These efforts culminated with the DOD Security Initiative of 1977. As a result of this initiative, audit capabilities were required in the DoD for resource-sharing systems containing sensitive information (Bace, 2000). The *ADP Security Manual*, articulated the following requirement in the early 1970's: "An audit log or file (manual, machine, or both) shall be maintained as a history of the use of the ADP [Automated Data Processing] System to permit a regular security review of system activity" (Department of Defense, 1973:38). While the previous example shows one instance of audit capabilities being mandated in DoD computer systems, more sensitive environments placed greater demands on the audit capabilities of resource-sharing systems (National Computer Security Center, 1985).

MULTICS was the most secure operating system compared to its contemporaries and offered the greatest promise in supporting multilevel secure computing environment (Karger and Schell, 2002). The changes in audit requirements for the MULTICS operating system illustrate a general trend in the use of audit functions as a security control. According to Karger and Schell, MULTICS audit capability as of the release of the 1972 Security Evaluation was a careful audit of the logins/logouts of each user (1974). Security enhancements proposed in 1973 at the Air Force Data Services Center expanded this audit capability to include:

- Accesses to classified data and the nature of the access (per DoD 5200.28-M)
- Each login and logout
- Unsuccessful login attempts and reason for rejection
- Rejected accesses of information based on security restrictions and each illegal attempted use (fault) of access permission
- All system faults which could indicate attempts to subvert the system or to exploit hardware failures
- All security-related actions of the SSO or the SA
- Each time a process awards itself extra privileges
- All completed requests for printed or punched output
- All tape mount requests for user tapes (Whitmore et al., 1973:80)

In 1975, Neumann et al. proposed a secure operating system design which included audit capabilities in its standard functions (1975). Still, a balance had to be struck between collecting too much audit data or not enough (Myers, 1980).

Intrusion Detection Is Born

In 1980, the Air Force again commissioned James Anderson. This time he was consulted to recommend improvements for the Air Force's security audit capability at a particular site. During this time, audit data was being manually analyzed. With an

increase in usage of computing equipment, came a corresponding increase in audit data. Anderson found that this volume increase coupled with an incomplete and sometimes redundant audit trail data, led to excessive time spent on auditing activities.

In laying a foundation for his discussion of audit trails, Anderson first provided a lexicon for discussing what he called a “security monitoring surveillance system” (Anderson, 1980: 4).

Threat: The potential possibility of a deliberate unauthorized attempt to:

- a) access information
- b) manipulate information
- c) render a system unreliable or unusable

Risk: Accidental and unpredictable exposure of information, or violation of operation integrity due to malfunction of hardware or incomplete or incorrect software design

Vulnerability: A known or suspected flaw in the hardware or software design or operation of a system that exposes the system to penetration or its information to accidental disclosure

Attack: A specific formulation or execution of a plan to carry out a threat

Penetration: A successful attack; the ability to obtain (unauthorized) access to files and programs or the control state of a computer system (Anderson, 1980:4-5)

Once a lexicon for investigating was introduced, Anderson pursued clarifying the “What should be detected?” question. Anderson’s report proposed a *threat taxonomy* to understand the types of threats and attacks that could be mounted against a computer system and how these threats may manifest themselves in audit data. Anderson identified three general cases of threats based on the attacker’s authority within the realm of the computer system (see Figure 2).

In the representation, Anderson noted that attackers are classified not with respect to the organization owning the system but with their authorization on the computer system itself. So, an external attacker may either be associated with the organization but

not authorized use of the computer system or the attacker could be a total outsider. In today's context of networks, this description could be extended to include individuals with access to the network, but not the target of the attack. Anderson also noted that the first task of an external penetrator is, in fact, gaining access to the system. Once access is gained, the threat is translated to an internal threat classification (Anderson, 1980).

With respect to individuals authorized on the specific computer system (i.e. internal penetrators), Anderson further identifies attackers as masqueraders, legitimate users, or clandestine users.

Masqueraders are those attackers who assume the identity of another user in order to carry out attacks. It's interesting to note that these users can be external penetrators who have assumed an identity internal to the system, or a legitimate user who has assumed an identity of another with malicious intent. Anderson writes that the masquerader is "interesting because there is no particular feature to distinguish the masquerader from the legitimate user" (Anderson, 1980:11). From the system's perspective, the masquerader *is* a legitimate user. As such, the masquerader is an "extra" use of the system by the unauthorized user" (Anderson, 1980:12). Hence, audit trail records should be able to detect masqueraders by assimilating the following indices:

- Use outside of normal time
- Abnormal frequency of use
- Abnormal volume of data reference
- Abnormal patterns of reference to programs or data (Anderson, 1980:12)

In summary, masqueraders are identifiable by distinguishing abnormal usage of a system.

This requires an assumption that those charged with security can discern normal usage.

	Penetrator Not Authorized to Use Data/Program Resource	Penetrator Authorized to Use Data/Program Resource
Penetrator Not Authorized Use of Computer	Case A: External Penetration	
Penetrator Authorized Use of Computer	Case B: Internal Penetration	Case C: Misfeasance

Figure 2: General Cases of Threats (Anderson, 1980:7)

Legitimate users are those individuals who have authorized access to a system and its resources but who misuse that access for whatever purpose. The degree of difficulty in detecting misfeasance by a legitimate user is greater in that no abnormal usage patterns may be present. The only way to identify the legitimate user is if the attack involves accessing information that is normally not authorized, accessing large amounts of information, or other relatively excessive usage patterns (Anderson, 1980). Again, detecting misfeasance by a legitimate user requires a previously established normal pattern of usage.

The clandestine user as characterized by Anderson is probably the most difficult to detect because this user is capable of hiding malicious activities by gaining the greatest

level of control over the system. Anderson felt that the ideal situation would be to provide independent audit trails for each major component of the system (1980). It is much more complicated to suppress auditing at multiple nodes in a networked environment. Conversely, analysis of audit trails from multiple systems is more difficult for security personnel (Anderson, 1980).

In addition to classifying internal intruders, Anderson also provided a framework for characterizing computer usage based on the parameters of time and dataset and program usage for individual users. Group statistics were used as a parameter for those systems requiring monitoring of particularly sensitive files and devices. Anderson proposed that correlating data between users and as a group would provide more granularity for what constitutes abnormal behavior. Anderson also proposed incorporating these data reduction techniques into an automated system which would alert system security officers when clear violation of system security policy or abnormal activity occurred (1980). This report served as a platform for the first IDS prototypes developed at SRI International and TRW.

An Intrusion Detection Model

From 1984 to 1986, Dorothy Denning and Peter Neumann, while employed with SRI, modeled a real-time intrusion detection system, called the Intrusion Detection Expert System (IDES), funded by the US Navy's Space and Naval Warfare Systems Command (SPAWAR). Taking the next step with Anderson's ideas, this model correlated anomalous activity with misuse using statistical analysis techniques. As such,

Denning's 1987 paper, "An Intrusion Detection Model", is one of the seminal works in the IDS discipline (Bace, 2000).

In Denning's paper, she describes the model as "being based on the hypothesis that exploitation of a system's vulnerabilities involves abnormal use, of the system; therefore, security violations could be detected from abnormal patterns of system usage" (1987: 1). Denning's model contains six components: subjects, objects, audit records, profiles, anomaly records, and activity rules. The model is based on creating *profiles* of normal behavior using statistical measurements to express the interaction between system subjects (users) and system objects (files, commands, devices, etc.). The profiles are compared to *audit records* to identify anomalous behavior and update the user's profile if necessary. When anomalous activity is identified, *anomaly records* are created. The model also includes *activity rules* that dictate the actions of the system in response to a stimulus, such as triggering a system alarm or generating a report when some predefined condition is satisfied. Denning considered the possibility that user behavior might change over time and included both static (long term measurements) and dynamic (short intervals) characteristics in the structure of profiles (1987).

Furthermore, the profiling mechanism had the capability of aggregating individual user profiles to create classes of users. This enabled the discovery of users whose behavior was internally consistent, but abnormal with respect to users with similar duties on the target system (Bace, 2000).

Between 1986 and 1992, Denning and Neumann's model was implemented in the IDES prototype at SRI International. One key feature of the IDES prototype was the

capability to alter a user's profile over time, based on changing usage patterns. While the ability to change a user's profile over time minimized the number of false alarms generated by the system, it also introduced the risk that a skilled hacker may gradually condition his own profile to defeat the system. In order to diminish this risk, a hybrid analysis scheme was adopted by the IDES development team which included an *anomaly detector* and an *expert system*. The anomaly detector employed the previously described statistical techniques to identify abnormal usage patterns. The expert system utilized a rule-based methodology to detect previously discovered attack patterns (Bace, 2000).

Henceforth, a dichotomy would emerge in the discipline of intrusion detection. *Anomaly detection*, as it has come to be known, involves the aforementioned statistical profiling techniques to discern abnormal user behavior from normal behavior. *Signature detection* (also called *misuse detection*) involves pattern-matching operations that search the audit logs for entries matching a predefined description of intrusive activity. The IDES prototype spurred a myriad of research IDS systems in the next decade, but most of them would employ only one of the two analysis schemes introduced in the IDES prototype (Bace, 2000).

IDS in the late 1980's

Many IDS systems were developed in the late 1980's subsequent to the landmark research completed at SRI. Each was significant in advancing the field of intrusion detection in its own way. Following is a sample of those systems along with notable accomplishments associated with each beginning with the Discovery system developed at TRW, the forerunner to the Experian credit reporting agency.

Discovery was an expert system for detecting problems in TRW's online credit database. This system was notable primarily because it could be considered the first application-based IDS system. Discovery was different from its contemporaries in that it monitored a database application for malicious behavior rather than an operating system. The purpose of the system was based on identifying three classes of abuse: unauthorized access, insider misuse, and invalid transactions. Discovery was funded and designed internally at TRW (Bace, 2000).

Haystack was a system developed for the US Air Force's Cryptologic Support Center initially by Tracor Applied Sciences, Inc. from 1987 to 1989 which was renamed Haystack Labs and continued the project from 1989 to 1991. The goal for Haystack was to aid security officers in detecting misuse of the Air Force's Standard Base Level Computers (SBLC). The SBLC were mainframes running early 1970's vintage operating systems and used for a myriad of tasks such as accounting, finance, inventory control, and personnel functions. The data residing in these systems was considered "sensitive but unclassified." These systems generated audit trails for over a million events per week (Bace, 2000).

The Haystack system was notable because it utilized an Oracle database management system (DBMS) to organize the audit trail data prior to analysis, as did later versions of IDES. Haystack performed audit trail analysis in batch mode. This meant that the system downloaded and analyzed audit data from the target system at regular intervals. As such, Haystack did not perform real-time intrusion detection, but was significant as it did monitor several geographically separated systems (Bace, 2000).

According to Innella, Haystack performed its analysis by comparing audit data with predefined patterns (2001). Axelsson also supports this, but added that Haystack incorporated both signature and anomaly detection in its analysis (2000a). The system earned its name purportedly because searching through the enormous volume of audit data for evidence of intrusion was like “looking for a needle in a haystack” (Innella, 2001).

The next important IDS development was the Multics Intrusion Detection and Alerting System (MIDAS). MIDAS monitored the National Computer Security Center’s Dockmaster system. Dockmaster was a resource facility for computer security that serviced the NSA, its vendors, academia, and other government agencies. Dockmaster provided e-mail and discussion services for its users using the early infrastructure of today’s Internet (Axelsson, 2000a). While similar to other contemporary systems in its employment of a hybrid analysis mechanism, MIDAS was important due to its status as the first IDS responsible for an operational system connected to the Internet (Bace, 2000).

The next major evolution in the intrusion detection field was the use of nonparametric statistical techniques for anomaly detection. This approach was implemented in a system called Wisdom and Sense by the Safeguards and Security Group at Los Alamos National Laboratory. Wisdom and Sense was developed to provide security for U.S. Department of Energy mainframes in several facilities. The use of nonparametric statistical mechanisms is significant because these techniques make no assumptions about the distribution of the data. In Wisdom and Sense, the nonparametric analyses were performed on archived audit data to formulate the rule base that

characterized normal activity. New activity was then compared to the rule base and deviations were catalogued. Another important feature of Wisdom and Sense is the capability for system administrators to augment the automated rule generation with tailored rules based on their specific knowledge of the system and its vulnerabilities (Bace, 2000). Wisdom and Sense suffered from many of the same shortcomings as other anomaly detectors of the era which included identifying appropriate data during the learning phase which was known to contain no intrusion activity, high false positives, and system memory limitations were insufficient to support expansive rule bases (Vaccaro & Liepins, 1989).

Up to this point, only host-based IDS systems had been developed, that is, they analyze data internally generated by a single system. Even in the case of Haystack, where several similar systems were being monitored, the detector proper performed analysis for only one of the systems at a time. The data was not correlated across the target systems. The next notable accomplishment in the evolution of ID systems is the Network System Monitor (NSM) at the University of California at Davis's Lawrence Livermore Labs. The emergence of the NSM marked the introduction of network traffic monitoring as the primary data source for intrusion data. All previous attempts at intrusion detection analyzed audit trails generated by the host operating system's internal security mechanisms or by keystroke monitoring mechanisms (Bace, 2000).

NSM accomplished network traffic monitoring by placing network interface hardware into *promiscuous mode*. This special operating mode allows the associated interface hardware to capture all the data traversing the transmission medium regardless

of the actual destination. By focusing on the network traffic, rather than internally generated audit trails, NSM was the first IDS that could be implemented in a heterogeneous computing environment (Bace, 2000). In other words, a single NSM platform could be employed to monitor systems running MS-DOS, Unix, or any variety of operating system so long as the system employed supported communication protocols and services, such as TCP/IP.

The NSM system was subjected to a rigorous two-month test where it monitored more than 110,000 connections on the Livermore Labs local area network (LAN) and appropriately identified 300 of those connections as intrusive activity. While NSM ascertained over 300 specific instances of misfeasance, the system administrators discovered less than one percent of abuses employing traditional manual methods (Bace, 2000).

The success of the NSM ultimately altered the IDS paradigm from host-generated audit trails to network traffic analysis. Most commercial systems used today rely on network traffic as their primary, if not sole, source of data. On a similar note, however, the creators of the NSM pointed out that a holistic approach to computer security would employ both network and host-based intrusion detection strategies (Heberlein et al., 1990). Furthermore, the developers of the NSM, employed both anomaly detection and misuse detection techniques in their analysis mechanisms, a trend which has continued throughout the evolution of IDS's to some degree.

As alluded to previously, the NSM marked the expansion of the IDS field from the host environment to the network environment. Also during this time, the internet

began experiencing significant increases in interconnectivity and bandwidth. Due to the exploitation of vulnerabilities such as the Morris worm of 1988, there arose an equally drastic increase in concern over computer security and an increase in research and development funding in both the academic and commercial environments. These exploits also brought about the creation of the Computer Emergency Response Team (CERT) by DARPA (McHugh, 2001). This brings us to the next major evolution of intrusion detection, the Distributed Intrusion Detection System (DIDS).

The DIDS initiative was the first integration of host and network-based intrusion detection capabilities. The importance of DIDS is reflected in the large-scale support and funding it received from three government agencies: the U.S. Air Force, the National Security Agency, and the U.S. Department of Energy (Bace, 2000). The development team was a virtual “Who’s Who in Intrusion Detection?” The team was comprised of many of the same minds who developed the NSM and Haystack.

Initially, DIDS was conceptualized as an aggregation of the techniques employed in the NSM and Haystack to detect intrusive behavior. However, the principal feature was to provide a centralized control and reporting console for system administrators. In order to accomplish this, DIDS overcame significant challenges. The first of which is tracking users and files in a networked environment. Intruders often exploit the interconnectedness of other systems as a platform for masquerade. In order to combat this, DIDS became the first system able to track users and objects (e.g., files) in the context of a networked environment. Another issue overcome by DIDS was that of correlating events from different layers of abstraction within a system. The designers

utilized a six-layer model where each layer was the result of a transformation applied to the raw data (Snapp et al., 1992). By overcoming these two key issues, DIDS became the first integrated tool for collecting and correlating evidentiary data relating to computer misuse (Bace, 2000). This is a key feature of today's IDS systems as they are often used as forensic tools for criminal investigation of computer crime.

The late 80's and early 90's also witnessed the birth of intrusion detection as a commercial product. While the systems themselves were not markedly different from those being prototyped in the research environment, it is noteworthy to chronicle the emergence of a market for IDS products. Three notable commercial products introduced during this period were ComputerWatch by AT&T, the Information Security Officer's Assistant (ISOA) by PRC, Inc. and Audit developed by Clyde Digital. ComputerWatch was an analysis tool for audit trail data with limited intrusion detection capability. AT&T made the product commercially available for a short time and then ComputerWatch was restricted to an internal resource for AT&T's consulting services group. ISOA was implemented in the UNIX environment and provided automated and interactive audit trail reduction and analysis. Audit scanned audit trails from Digital Equipment Corporation's VAX/VMS operating systems for anomalous activity and ranked and reported users based on their level of suspected misuse (Bace, 2000). It is interesting to note that Clyde Digital eventually became Axent, which is now owned by Symantec. Symantec entered into the commercial IDS market by acquiring the Intruder Alert (host-based) and NetProwler (network-based) technologies from Axent (Innella, 2001).

Surveys of IDS Research

Since the birth of intrusion detection as a viable commercial market in the mid-1990's, IDS has grown to a major information security market, amassing \$382 million in worldwide annual revenue in 2002, according to a press release from Infonetics Research (2003). This same report extrapolated market growth to the \$1.6 billion level by 2006.

In spite of the apparent commercial success of intrusion detection systems in recent years, both researchers and end users of IDS technology have identified gaps in the current body of research which manifest themselves as shortfalls in operational systems. These shortfalls are most evident in the high false alarm rates of ID systems and the proprietary designs of many ID systems that fail miserably in their interoperability with other ID systems and security products.

A few researchers have attempted to survey the IDS literature and provide guidance for further development of the technology. Lunt described the prototype systems at the time which included IDES, MIDAS, Discovery and others as well as summarizing their respective detection approaches. She concluded the survey with the insight that a successful IDS “should incorporate several different approaches” and implied that IDS security was of some concern (Lunt, 1988:15).

In 2000, Allen et al. also surveyed the ID field and formulated a series of challenges to current ID systems. These challenges can be summarized as follows:

- Increase in number, variety, and sophistication of attacks
- Ubiquity of strong encryption which hinders network IDS's
- Need for interoperability and correlation of security data from heterogeneous infrastructure
- Ever-increasing network traffic with respect to volume and speed, coupled with proliferation of network switching systems

- Lack of widely accepted terminology and foundational concepts in ID
- Security of intrusion detection systems from tampering
- Poor security design of common operating environments
- Absence of formal procedures and testbeds for objective evaluation of ID systems
- Unacceptably high levels of false positives and false negatives (Allen et al., 2000)

The report went on to make recommendations to sponsors, users, vendors and researchers to address the gamut of issues facing the IDS community.

One of the more compelling recommendations included a call to both vendors and researchers to focus more energy into cross-platform correlation of data coupled with diverse detection approaches, or data fusion. Most current commercial systems use primarily misuse detection approaches (i.e. pattern matching). Others support both network and host-based intrusion detection, or hybrid IDS, but fail to correlate the diagnoses from the two components. Allen et al. felt that research into data fusion would significantly reduce high false alarm rates that are common in current IDS systems (2000).

Along the same lines, the authors felt that research efforts in ID should focus on the foundational questions in the discipline. What should be detected? How should we detect it? To quote them,

“We believe this should be a wake-up call to the research community. Research efforts in the intrusion detection field have focused increasingly on building ID frameworks that address higher level issues and abstractions. We believe that these efforts should, in great part, be redirected to address the fundamental issues of minimizing false alarms and quantitative testing of the resulting algorithms. If the false alarm issue is not resolved, all the frameworks in the world will not help. (Allen et al., 2000:79)”

Another major theme was that of evaluating IDS systems and signatures. Some problems found in evaluations of commercial products were of such a basic nature that

there was a question whether the vendor performed any in-depth testing of the system at all. The implication is that vendors are riding the IDS bandwagon to success with little justification for the advertised effectiveness of their products. The authors recommended that the IDS research community explore methodologies for testing IDS systems and also proposed that vendors publicly release their attack signatures to scrutiny, much like the anti-viral community which has flourished under this model. Thus, the signatures created would be quickly disseminated as well as robust to new attacks soon after they appear (Allen et al., 2000).

Sill another recommendation expanded on the idea of utilizing IDS systems in gathering and collecting forensic evidence. Allen et al. stated that IDS technologies possess vast potential in forensic evidence collection as they are capable of “time-lining suspicious activities such as network scans, login attempts and document modifications” (2000:82). Currently, the false alarm problem remains a confounding variable to employing IDS in this fashion. While IDS remains useful as a platform for investigation, maintaining the chain-of custody and insuring against evidence tampering require the use of traditional manual evidence collection processes in order to provide legally admissible evidence.

McHugh echoed many of these very same sentiments in his brief survey of the field in a 2001 analysis of the intrusion detection discipline. He concluded that the central issues in resolving detection accuracy rested on two assumptions made during early work in the field that continue to manifest themselves in current approaches. The first is that intrusions would manifest themselves in such a way that the presence of an

attack is obvious and could be codified in such a way to allow for easy detection. The other assumption was that any activity outside of “normal” activity is a strong indicator of intrusive behavior (McHugh, 2001).

McHugh proposed two lines of research that he believed would help in clearing up these foundational issues in the ID field. First, he proposed research in characterizing “normal” behavior. By modeling normal behavior, it would make designing anomaly detectors much easier since they are based on identifying activity that is outside of declared or learned norms for computing behavior. The second proposed line of research, strongly related to the first, involves formulating a credible theory for intrusive behavior (McHugh, 2001).

These proposals borrow from the general signal detection theory (see Figure 3). A study by Axelsson explored the connective threads between the classical detection theory and intrusion detection. In classical detection theory, two signals are present. In order to carry out detection, the detector should be able to distinguish both signals in order to accurately carry out its objective. In the context of intrusion detection and currently employed detection approaches, the detectors operate with information on only one of the signals. In the anomaly detection approach, at the most fundamental level, the intrusion detection system possesses knowledge of the “noise” and detects that which is not noise. Similarly, in misuse detection, the intrusion detection system possesses partial knowledge of the signal (since not all possible attacks are known) and reacts when that signal is detected. He concluded, “If we wish to classify our source behaviour

correctly...knowledge of both distributions of behaviour will help us greatly when making the intrusion detection decision” (Axelsson, 2000b:9). It is interesting to




Figure 3: Classical Detection Theory (Axelsson, 2000b)

note that in either approach, the detector operates with only partial knowledge of the signal it's trying to detect. This is because research has not sufficiently defined normal behavior nor intrusive behavior to a sufficient degree to alleviate the false alarm problem.

Axelsson also produced a survey and taxonomy of intrusion detection systems which noted interesting trends in both research and commercial ID systems being developed. The author noted that current trends in the discipline included active response, IDS security, and interoperability among intrusion detection systems, trends also alluded to by Allen et al. (2000). In addition, Axelsson identified a general trend toward distributed intrusion detection systems rather than a centralized management and

analysis platform, adding that this trend followed a more general trend in the computing environment (2000a).

Lundin and Jonsson produced a similar survey in 2002. The authors specified four areas as the most important to the advancement of the field. Those areas were:

- Classifying intrusions - without a tried and true methodology for classifying intrusions or classes of intrusions, it is impossible to determine what data is needed to detect them
- Effective data collection - the research community must define in detail what data must be collected to allow detection and devise methodologies for collecting the data efficiently
- Data visualization (a.k.a. alarm reporting) - security personnel must have an acceptable interface for viewing and exploring the intrusion data
- Methodology for generating test data - must be easier and automated to allow for robust evaluations of new systems (Lundin & Jonsson, 2002:34-35)

They arrived at the conclusion that most effort was being placed into agent-based intrusion detection and the development of IDS evaluation testbeds. Also mentioned in the report was the increased interest in shoring up security for intrusion detection systems and to active response capabilities (Lundin & Jonsson, 2002). Still, while the research community drives some aspects of the intrusion detection discipline, the commercial market also helps to drive advances in the field based on customer desires.

IDS Market Trends

Recent forecasts for the IDS market show a significant interest in a new generation of intrusion technologies coined intrusion detection and prevention systems (IDP), or intrusion prevention systems (IPS). The feature that overtly distinguishes intrusion prevention is the capability to automate responses to intrusive behavior.

Possible scenarios include simply stopping certain types of traffic much like a firewall, or redirecting the traffic to a safe destination for further surveillance.

Some of the top vendors have begun deploying solutions that consolidate and integrate many of the major security technologies into a single appliance that can be strategically inserted into existing network architectures. These appliances typically have consolidated IDS/IPS technology with firewall, antivirus, and/or vulnerability assessment technologies. Figure 4 illustrates the recent flattening in revenues for traditional IDS products while in-line IDS appliances are on the rise and expected to surpass the traditional implementations in the future.




Figure 4: Worldwide IDS/IPS Product Revenue Breakdown by Year (Infonetics Research, 2004)

Another current trend in security is that of utilizing the packet inspection capabilities of IDSs as a policy enforcement tool for corporate firewalls. This implementation places the sensor beyond the “safe” region, or DMZ, which is the buffer zone between the firewall and the Internet. By placing the IDS outside the firewall, analysts are able to ensure that firewalls are, in fact, configured according to the corporate security policy. Since numerous manifestations of malicious code utilize common protocols such as HTTP and FTP to defeat firewall rules, IDS systems can identify many of these attacks and trigger a quicker response from administrators.

By examining the historical background surrounding the emergence and evolution of the intrusion detection discipline, it’s clear that the research covers a broad range of topics and issues within the realm of computer security. It is also clear that while this huge body of research exists, the application is lacking due to significant gaps in the foundations of the discipline. However, what lies at the heart of this research is whether IDS has a place within an information assurance program. In order to reach any reasonable conclusion regarding the role of IDS in information assurance, an examination of the information assurance literature is necessary to set up a framework for addressing intrusion detection in the context of IA.

Information Assurance

One of the earliest appearances of the term “information assurance” was in Joint Publication 3-13 *Joint Doctrine for Information Operations* published by the Department of Defense to define concepts for utilizing information in the military environment. It was defined in this publication as those measures

“that protect and defend information systems by ensuring their availability, integrity, authentication, confidentiality, and non-repudiation. This includes providing for restoration of information systems by incorporating protection, detection, and reaction capabilities” (Department of Defense, 1998:I-9).

IA in the DoD, is considered a subset of defensive information operations and according to Air Force doctrine encompasses what was once known as INFOSEC (information security) (Department of the Air Force, 2002). As many military concepts are adapted to meet private sector needs, the concept of IA also followed this model and is now a widely accepted construct in both private and public sector organizations.




Figure 5: Information Assurance (Department of Defense, 1998)

As evidenced by the preceding definition, there are five elements to IA: availability, integrity, authentication, confidentiality, and non-repudiation (see Figure 4).

An understanding of each of the elements is necessary to appreciate the scope of the IA construct. The five elements of IA could be described as follows:

Availability: information can be accessed when and where required and in the proper format

Integrity: information is accurate and timely

Authentication: information is available only to those who have proof of authorization

Confidentiality: information is only available to those who require it

Non-repudiation: information cannot be forged nor can its origin be denied
(Maconachy et al., 2001; McKnight, 2002)

The Department of Defense operationalized the IA construct in its formulation of the *defense-in-depth* strategy which has also been widely adopted in the private sector. Defense-in-depth takes a truly strategic view of IA in that it advocates a balance between capability, cost, performance and operational considerations (National Security Agency, 2002). A key principle of defense-in-depth is that IA can only be achieved through the interaction of people, technology, and operations.

While people and operations are clearly important elements in an IA strategy, this research deals primarily with the employment of technology to fulfill security requirements. The following illustration describes the interaction between the three elements, but also specifies four focus areas that encompass the primary objectives for technology in a defense-in-depth strategy: defend the network and infrastructure, defend the enclave boundary, defend the computing environment, and supporting infrastructures.

Defense-in-depth also specifies five principles that should be followed in order to achieve IA: defense in multiple places, layered defenses, security robustness, deploy KMI/PKI (key management infrastructure/public key infrastructure), and deploy intrusion detection systems. By spreading out defense mechanisms throughout the




Figure 6: Defense in Depth Strategy Focus Areas (National Security Agency, 2002: 2-11)

infrastructure, organizations are better able to respond to attacks, whether they originate from inside or outside the network, or originate from single or multiple sources. By layering defenses, an organization capitalizes on the strengths of multiple protection mechanisms, rather than risking damage due to the breach of a single line of defense. Security robustness refers to measuring and reporting the level of protection of a specific IA component based on its strength to weather attacks and the assurance that it is deployed and operated properly. The last two principles recommend employing specific technologies including IDS and KMI/PKI (National Security Agency, 2002). The researcher does not consider these as elements of a strategy, but rather an implementation plan.

Summary

This chapter examined the roots of the intrusion detection discipline within the context of the more general computer security discipline. Chapter Two also traced the evolution of intrusion detection from its inception through current research in the field. Finally, Chapter Two investigated the concept of information assurance along with its associated constructs in order to establish a framework for discussing intrusion detection within the context of IA strategy. In the following chapter, the historical analysis methodology is explored.

III. Methodology

Chapter Overview

Chapter Three explores the historical analysis methodology and the inherent value of historical research. This chapter will present various constructs relating to historiography proposed by previous researchers in MIS as well as general management and business. Furthermore, the researcher will provide justification for employing this methodology in the context of the problem proposed in Chapter One.

The Importance of History

“History is the witness that testifies to the passing of time: it illumines reality, vitalizes memory provides guidance in daily life and brings us tidings of antiquity.” (Cicero)

Despite advocacy by great minds such as Cicero, the historical perspective is one that has been sparsely employed in information systems (IS) research. Bannister wrote that rigorous historiographies in IS research are scarce (2002). Still, historical perspectives can provide significant insight into contemporary situations.

According to Leedy and Ormrod, history itself is nothing more than a collection of events and changes in the human condition (2001). Historical research, however, tries to make sense of these events and changes and assign to them a greater meaning (Leedy and Ormrod, 2001). This process fosters edification. According to Mason et al., the great economist Joseph Schumpeter asserted that for any field of study to be considered a discipline, it must provide “(1) empirical data, observations and facts, (2) theories and paradigms, (3) ethics, and (4) history” (1997a:257-258). The authors expanded, “history

is necessary to provide a temporal and contextual meaning for each of the other three forms of knowledge” (Mason et al., 1997a:258). Parallel to this concept, O’Brien et al., stated that knowledge is only effective when it is contextualized, and that history is the vehicle by which we obtain contextualization (2004). Ergo, the value of the historical method is in the power of the past to enlighten the present.

The Historical Method

As the value of history has been unequivocally stated, it is also necessary to explore the discipline-specific approaches to historiography in IS research. The existing approaches can best be explained by a discussion of the various works and their associated constructs relating to the historical method.

The approach proposed by Mason et al. incorporated two constructs that could be used to investigate the evolution of IS within an organization. The first construct dealt with the roles of specific individuals in affecting the organizational evolution and attainment of a *dominant design*, a new configuration of technology, strategy, and structure which gives the organization a competitive advantage. As this research does not focus on individuals per se’, this construct will not be employed during the course of the analysis. The second proposed construct, the *cascade*, describes five evolutionary phases of an organization following the recognition of a crisis (Mason et al., 1997a). These five phases are illustrated in Figure 6. Bannister, however, noted that this approach is constrained by the assumption of a major crisis provoking the organization to action. He argued that not all organizations achieve dominant design after negotiating major crises.




Figure 7: The Cascade (Mason, McKenney, & Copeland, 1997a:267)

Rather, Bannister followed Greiner's rationale that organizations continuously tackle crises on a regular basis (2002). This researcher believes that these crises, or opportunities depending on your perspective, prompt organizations to continually adjust their dominant design to meet the needs of their environment.

In a supporting work, Mason et al. fielded a seven-step process for engaging the historical method. The seven steps proposed are:

- 1) Begin with focusing questions—specify those related areas to be addressed in the course of research
- 2) Specify the domain—determine the range of events that pertain to the research question(s)
- 3) Gather evidence—provides material for historical interpretation and contextualization
- 4) Critique the evidence—sources are screened for their reliability and relevance
- 5) Determine patterns—emergent patterns and theories are codified
- 6) Tell the story—the account; present the evidence in a manner consistent with reality
- 7) Write the transcript—conclusions are placed in context with existing schemas (Mason et al., 1997b:312-317)

Mason attested that these steps do not require stringent adherence and the phases often overlap (1997b). Such is the nature of historical research, a dynamic process that requires many revisits to properly contextualize the topic in question.

O'Brien et al. proposed a similar framework for use in business and management studies. This framework, while comparable to that proposed by Mason, contains a nine-step process outlined below:

- 1) Develop research question—ensure that the topic is robust to historiography, interesting, defensible, feasible

- 2) Relevance check—researcher must ensure usefulness of historical perspective
- 3) Scope of the research—define the domain associated with the focusing question(s)
- 4) Sources of evidence—identify valuable/reliable data sources
- 5) Assess methods of analysis—researcher must decide on qualitative, quantitative, or hybrid approach to analysis
- 6) Assemble the evidence—researcher must sift sources to ensure reliability/accuracy
- 7) Develop the story—requires identification of emergent patterns and interpretation of facts; answer the what, how , and why
- 8) Critique the story—ensure that the evidence has been critically evaluated and presents a persuasive argument
- 9) Outcome of the research—communicate the conclusions of the research, including limitations and implications (O’Brien et al., 2004: 138-141)

On close examination of each of these frameworks, convergence between the two authors is transparent. While the two approaches are not completely synchronous, each is robust to addressing the research at hand and provides a foundational framework from which to engage the problems proposed in chapter one. The Mason framework provides value in that it espouses an organizational perspective for understanding how IT is introduced and the factors that contribute to or constrain its usefulness. The O’Brien framework, on the other hand, does not limit its usefulness specifically to IT as it applies to general business and management. Due to the focus of this research, that of providing a strategic understanding of a particular technology for use by executives in making business and management decisions, it seemed intuitive that a synthesis of the two frameworks would be useful and applicable to this study. The researcher identified

parallel concepts between the two constructs in order to formulate the synthesis, which is illustrated in Figure 8.

Mason et al. (1997b)	Mason-O'Brien Synthesis	O'Brien et al. (2004)
Begin with focusing questions	Develop focusing questions	Develop research question Relevance check
Specify the domain	Determine scope	Scope of the research
Gather evidence	Gather and review evidence	Sources of evidence Assess methods of analysis
Critique the evidence	Critique evidence	Assemble the evidence
	Assemble evidence	
Determine patterns	Identify emergent patterns	Develop the story Critique the story
Tell the story Write the transcript	Convey results	Outcome of the research

Figure 8: Mason-O'Brien Synthesis

Justifying an Historical Analysis

As alluded to previously, historical research provides a valuable lens through which to perceive substantive phenomena. Knowledge and comprehension of past events sheds light on present circumstances and enables the creation of mental models for future situations. In the context of information systems, historiography is capable of illuminating those circumstances that affected the introduction of technology to address practical issues.

While historiographies are relatively underutilized in IS research, they play a valuable role regardless. Mason, McKenney, & Copeland stated, “Historical analyses broaden our understanding of the processes by which information technology is introduced into organizations and of the forces that shape its use” (1997a:257).

Furthermore, the authors identified three specific products of historical analyses that are of analytical value in IS research:

- 1) First is an account of a significant fragment of the past describing events of importance to the MIS community which also provides contextual material for understanding other events
- 2) Second, the historical account can be used as data in a broad process of inductive reasoning
- 3) Third, the research may serve as the source of new research hypotheses (Mason et al, 1997b:308-309)

Bannister concurred with Mason et al, but he added that historical research is also profitable for validation or falsification of existing theory (2002).

While Mason et al. proposed historical analysis as a valuable apparatus for understanding organizations, this methodology can also offer significant insight for understanding other entities. O'Brien et al. proposed that researchers may choose organizations, individuals, or industries as units of analysis during the scoping of the subject matter; hence, historical analysis is appropriately applied to the ID discipline in this research (2004).

Limitations of the Historical Method

As with any methodology, there are inherent constraints which the researcher must acknowledge and strive to minimize. First and foremost, historical research, according to O'Brien et al., is "highly interpretist and the findings are often thought to be more personal than some researchers are comfortable with" (2004:143). As such subjectivity is introduced, it must be foremost in the researcher's mind to focus as much energy as possible in interpreting factual information. Still, as Leedy and Ormrod point

out, it is nearly impossible for researchers to avoid working with data contaminated by bias. What is important, though, is that the possible presence of bias be acknowledged, and in the case of historiography, this is arguably a greater possibility than in other methodologies (Leedy and Ormrod, 2001).

A second limitation with the historical method is the possible incompleteness of historical data. Elton says that in employing this method, the researcher becomes “a servant of his evidence,” according to O’Brien et al. (2004:139). O’Brien also noted that more often than not, the historical researcher is not working from complete information. This implies that the researcher, in order to maintain the integrity of the inquiry, must continually refine his evidence according to an espoused sense of credibility and reliability. Furthermore, some evidence will meet the acceptance criteria for one researcher where another would discount the same source (O’Brien et al., 2004).

Summary

This chapter discussed the proposed historical research methodology. It discussed the importance of history in general as well as the value of historical research in particular to the IS discipline. Chapter Three outlined two approaches to historical research and noted their similarities and followed with an exploration of the proposed constructs corresponding to each approach. Furthermore, this chapter presented a justification for this approach within the context of the IS discipline and specifically within this research effort. Finally, the author identified possible limitations of the historical research method.

IV. Findings

Chapter Overview

Chapter Three described the historical research methodology and provided justification for its use in resolving the investigative questions posed in Chapter One. This chapter communicates the results of the historical analysis as applied to the factual data presented in the literature review.

Addressing the Investigative Questions

This section specifically addresses the three investigative questions posed in Chapter One. By answering these questions, IS/IT executives will be better able to understand the background of intrusion detection and its related technologies. Furthermore, by contextualizing the IDS discipline, IS/IT executives will be better informed to make strategic decisions regarding the employment of intrusion detection within their respective business domains.

The methodology employed to address these questions was based on a synergy of the historical approaches proposed by Mason, McKenney, & Copeland in 1997 and O'Brien, Remenyi and Keaney in 2004. In establishing the existence of the problem, the researcher developed the focusing questions and scoped the research primarily by becoming familiar with the subject matter and the established history of intrusion detection. This focusing of the research was followed by an initial review of the historical data and seminal works within intrusion detection. Following the initial

review, the researcher critiqued the evidence based on applicability and relevance to the focusing questions and led to the identification of additional sources. Then, the researcher assembled the relevant evidence to tell the historical account which subsequently allows emergent patterns to develop and be identified. The last step of the synergized methodology is to convey the results of the research. This will be accomplished by addressing the investigative questions with respect to the manner in which they were presented in Chapter One.

Investigative Question #1: What factors led to the emergence of intrusion detection as a discipline?

The first of the three investigative questions posed in response to the problems identified in Chapter One was that of identifying those factors which led to the emergence of the intrusion detection discipline. It stands to reason that any matter of inquiry that has effloresced into an academic discipline must serve a purpose and provide contribution to the general body of knowledge. In answering this question, the researcher hopes to shed light on the reasons for which this discipline appeared and provide clarity as to whether the discipline has or has not outlived its purpose as concluded by Gartner (2003).

By applying the cascade construct as a vectoring tool, the data identified three factors which contributed to the emergence of the intrusion detection discipline:

- I. Increased acquisition and usage of resource-sharing systems in the Department of Defense*
- II. Growing need to employ resource-sharing systems within an open computing environment while maintaining security*

III. Unmanageable volume of audit data being produced

These three factors interact to form the crisis that led to intrusion detection developing into its own academic discipline.

Investigative Question #2: How has IDS research evolved since its inception?

Now that the research has identified the circumstances underlying the emergence of the intrusion detection discipline, it is now necessary to explore how the discipline has evolved since its inception in order to adequately contextualize the topic under study.

Six trends were identified that could be used to describe the evolution of IDS research from its inception through current research:

- I. From passive to active response mechanisms*
- II. From centralized to distributed analysis*
- III. From centralized to distributed/agent-based collection*
- IV. From single to multiple detection approaches within a system*
- V. From host-based analysis to network-based analysis to hybrid analysis*
- VI. From software-based systems to hardware appliances/in-line devices*

These six trends provide granularity for comprehending the growth of the ID discipline in depth and breadth over time. In answering the first two investigative questions, the research has provided contextual data for understanding the inception and evolution of intrusion detection systems. The results of the first two investigative questions as well as the contextual data used to address them were then synthesized to address the final investigative question.

Investigative Question #3: What specific roles, if any, are IDS's suited to fulfill within a holistic information assurance program?

While it is clear that IDS on its own does not currently possess the capability to detect all known and unknown intrusions, this research shows that, when strategically employed, IDS can serve a significant purpose within the context of the information assurance construct. As far as the five pillars of IA, intrusion detection can be categorized as supporting, to some degree, all of the primary elements. For example, IDS supports confidentiality by reporting when a user has accessed certain privileges typically afforded to administrators or when those privileges have been employed beyond a predetermined normal threshold. This particular example also shows evidence of supporting integrity (e.g. IDS may aid in identifying resources that the attacker gained access to) and authentication (e.g. anomaly detectors allow profiles of usage to be created which can highlight intrusive behavior). Another example would be IDS supporting availability by indicating when a network is being targeted for denial-of-service attacks and non-repudiation by logging the origin of the attack. Clearly, due to the design constraints of intrusion detection systems, it may not be possible to provide support for all elements of the IA construct all the time. As such, the technology is properly designated a supporting infrastructure within the defense-in-depth strategy. However, the researcher does feel it is necessary to point out the contradictory nature of identifying specific technologies to be employed when articulating organizational strategy.

With respect to the defense-in-depth operational strategy, certain principles in IA clearly support overlap and redundancy in the deployment of protection mechanisms, specifically, defense in multiple places and layered defenses. No one security technology

is capable of effectively shielding an entire network from all possible attack scenarios. Therefore, IDS when employed strategically is just another tool in the toolbox to achieving information assurance. Furthermore, the IA construct appears to advocate IDS in its inclusion of “protection, detection and reaction capabilities” (Maconachy et al., 2001:307) which support system restoration following an attack. With this in mind, the research identified three primary roles that intrusion detection systems are capable of fulfilling within the context of information assurance.

- I. Intrusion detection can be useful as a stimulus to actuate a predefined response mechanism*
- II. IDS can be used to gather, organize, and correlate evidence in the investigation of computer misuse*
- III. IDS can be employed as a vulnerability assessment/policy enforcement tool*

The role that IDS is able to fulfill ultimately depends on the objectives of the organization employing the technology. One organization may wish to furnish forensic evidence to law enforcement officials in order to aid in prosecution of a computer crime while another may wish to simply block intrusions and prevent further damage. Furthermore, these roles are not mutually exclusive. If so inclined, it is possible for an organization to deploy IDS such that it fulfills all of these roles in concert.

A Synthesis

The overarching goal of this research effort was to provide IS/IT managers guidance regarding the utility of intrusion detection systems within the framework of

information assurance. The method chosen to accomplish this objective was historical analysis due to the contextualizing value that this methodology provides.

Past researchers in IDS have proposed that the commercial success of IDS technology has been contingent upon a bandwagon effect. Organizations implemented the technology without a strategic perspective on its capabilities and limitations. IDS investment decisions were based on novelty rather than feasibility. This has led to relative frustration with solutions that fail to deliver what executives thought would be delivered. In many cases, IDS was viewed as a total solution to information security problems.

The question of deploying IDS is not answered with a binary response. By providing insight on how the discipline began and describing the evolution of intrusion detection, this research gives managers a more realistic understanding of the capabilities and limitations of IDS technology. While this knowledge alone may have better prepared organizations to deal with the confounding issues in implementing IDS, this research further proposes specific roles that IDS technologies are suited to address within an IA strategy. This knowledge and understanding of IDS technologies prepares executives to make more informed decisions regarding the usefulness of IDS technologies within their own organizations.

Summary

Chapter Four outlined the results obtained by applying historical research techniques to the topic of intrusion detection. This chapter discussed the three investigative questions posed in Chapter One and provided conclusions based on an

application of the historical research methodology. The final chapter contains a summary of the research findings and proposes practical implications and recommendations for future research.

V. Summary, Implications, and Recommendations

Summary

This research investigated the historical background and circumstances that led to the birth of the intrusion detection field and explored the evolution of the discipline through current research in order to identify appropriate roles for IDS technology within an information assurance framework. The research identified factors contributing to the birth of ID including increased procurement and employment of resource-sharing computer systems in the DoD, a growing need to operate in an open computing environment while maintaining security and the unmanageable volume of audit data produced as a result of security requirements. The research also uncovered six trends that could be used to describe the evolution of the ID discipline encompassing passive to active response mechanisms, centralized to distributed management platforms, centralized to distributed/agent-based detection, single to multiple detection approaches within a system, host-based to network to hybrid analysis and software-based to hardware-based/in-line devices. Finally, the research outlined three non-mutually exclusive roles suitable for IDS to fulfill within the IA framework including employing IDS as a stimulus to incident response mechanisms, as a forensic tool for gathering evidence of computer misuse and as a vulnerability assessment or policy enforcement facility.

Practical Implications

The primary contribution of this research is in supplying IS/IT executives and Chief Information/Security Officers with a usable framework with which to make strategic decisions to deploy intrusion detection systems. It is beyond the scope of this research to determine whether or not IDS is appropriate for a given organization. This is a decision that must be appropriated to ranking executives who possess the corresponding strategic perspective required to assess the suitability of the technology. However, the research concluded that IDS is robust to addressing specific roles should corporate strategy dictate their relevance.

The historical analysis provides a distinct perspective that can be used to contextualize the current state of the intrusion detection discipline. Furthermore, it provides executives with an understanding of not only the limitations of the technology but how those limitations developed over time. This knowledge alone supplies managers with a powerful tool for gauging investments in current and evolved forms of the IDS paradigm such as intrusion prevention systems or intrusion detection and prevention solutions. Thus, this research effort may aid in formulating long-term strategies for employing IDS technologies or not employing them, depending on organizational objectives.

While the technology is a central focus of this research, it is important to note that a holistic information assurance strategy includes both personnel and organizational issues in addition to technology. Because current IDS technology is not yet able to completely replace the security analyst, it becomes paramount for executives to acquire

and train skilled personnel to interact with IDS systems as well as align organizational structures to garner the greatest effectiveness from IDS technologies. Thus, executives must continually weigh personnel issues in their strategic employment of IDS technology.

The researcher does believe, in contrast to Gartner, that accuracy in future IDS implementations will realize sufficient improvement as to allow more active response capabilities. Detection accuracy improvements will only be realized, however, when security technology vendors implement standards for the exchange of security information. By fusing security information from the entire gamut of security technologies, detection will occur quicker and with ever-increasing accuracy. As such, executives should actively pursue bolstering their infrastructures with interoperable technologies. By coupling interoperability with advanced visualization tools, security personnel will be postured to take calculated action to repel and recover from malicious attacks.

Recommendations for Future Research

As this was an exploratory study, and highly interpretist, a similar study on the same topic may yield different results. However, should another researcher, surveying a similar body of literature, arrive at comparable conclusions, greater validity could be attributed to this research effort. Furthermore, future research may identify other roles for which intrusion detection is appropriate within the IA framework of protective measures.

Future research could also focus on describing how current organizations actually employ IDS with respect to the roles identified in this inquiry. By identifying which roles IDS fulfills more often in practice, IT decision-makers could be better suited to weigh decisions on IDS investment. Along the same lines, individual implementations of IDS technology could be evaluated on their ability to perform within the context of the roles identified in this research.

Another beneficial future research effort would be a cross-disciplinary effort between IDS researchers and behavioral scientists. As this effort and others have shown, the detection problem is largely due to a lack of foundational research in modeling intrusive and normal behaviors during human-computer interaction. Intuitively, it would seem obvious to borrow from the behavioral sciences, which are much more adept at modeling human behavior, to address this gap in IDS research.

Another possible avenue for further investigation would be the effect of technology convergence within the framework of IA. As IDS, firewall, antivirus, and vulnerability assessment technologies converge, it would be fruitful to investigate the effects of this synergy on organizational IA strategies. The effect would be particularly interesting in light of the defense-in-depth principles of layering and distributing defenses throughout the network. Furthermore, some of the evidence seemed to present a trend toward enterprise security solutions provided by commercial ISP's for their customers. These services are often transparent to the user and could include intrusion detection, anti-virus, spam filtering, and spyware detection and removal services. An exploratory study into the extent and nature of this trend may yield beneficial results.

Bibliography

- The Worldwide Michelangelo Virus Scare of 1992*. Retrieved February 2, 2004 from http://www.vmyths.com/fas/fas_inc/inc1.cfm
- Allen, J., Christie, A., Fithen, W., McHugh, J., Pickel, J., & Stoner, E. (2000). *State of the Practice of Intrusion Detection Technologies*. Technical Report No. CMU/SEI-99TR-028. Carnegie Mellon University, Software Engineering Institute.
- Anderson, J. P. (1980). *Computer Threat Monitoring and Surveillance*. Fort Washington PA: James P. Anderson Co.
- (1972a). *Computer Security Technology Planning study, Volume I*. ESD-TR-73-51. Hanscom Field, Bedford MA: ESD/AFSC.
- (1972b). *Computer Security Technology Planning Study, Volume II*. ESD-TR-73-51. Hanscom Field, Bedford MA: ESD/AFSC.
- Axelsson, S. (2000a). *Intrusion Detection Systems: A Survey and Taxonomy*. Technical Report No. 99-15. Goteborg, Sweden. Chalmers Univ. of Technology.
- (2000b). *A Preliminary Attempt to Apply Detection and Estimation Theory to Intrusion Detection*. Technical Report. No. 00-4. Goteborg, Sweden: Chalmers Univ. of Technology.
- Bace, R. G. (2000). *Intrusion Detection*. Indianapolis: Macmillan Technical Publishing.
- Bannister, F. (2002). "The Dimension of Time: Historiography in Information Systems Research," *Electronic Journal of Business Research Methods*, 1(1), 1-10.
- Denning, D. E. (1987). "An Intrusion Detection Model," *IEEE Transactions on Software Engineering*, 13(2), 222-232.
- Department of Defense. (2002). *Information Assurance*. Department of Defense Directive 8500.1.

- , (1998). *Joint Doctrine for Information Operations*. Joint Publication 3-13.
- , (1973). *ADP Security Manual: Techniques and Procedures for Implementing, Deactivating, Testing and Evaluating Secure Resource-Sharing ADP Systems*. DOD 5200.28-M.
- Department of the Air Force. (2002). *Information Operations*. Air Force Doctrine Document 2-5.
- Gartner, Inc. (2003). *Gartner Information Security Hype Cycle Declares Intrusion Detection Systems a Market Failure*. Retrieved February 2, 2005 from http://www3.gartner.com/press_releases/pr11june2003c.html
- Heberlein, L. T., Dias, G. V., Levitt, K. N., Mukherjee, B., Wood, J., & Wolber, D. (1990). "A Network Security Monitor," *Proceedings of the 1990 IEEE Symposium on Research in Security and Privacy*, Oakland CA, 296-303.
- Infonetics Research. (2004). *ISS and Cisco Tie for Lead in IDS/IPS Market, Prevention Drives Market Growth*. Retrieved January 21, 2005 from <http://www.infonetics.com/resources/purple.shtml?ms04.id.3q.nr.shtml>
- , (2003). *Worldwide Intrusion Detection/Prevention Product Revenue Hits \$382M in CY02; Technology Transition Drives Market Explosion*. Retrieved January 21, 2005 from <http://www.infonetics.com/resources/purple.shtml?nr.idsms.4q02.030403.shtml>
- Innella, P. (2001). *The Evolution of Intrusion Detection Systems*. Retrieved January 12, 2005 from <http://www.securityfocus.com/infocus/1514>
- Innella, P., & McMillan, O. (2001). *An Introduction to Intrusion Detection Systems*. Retrieved February 2, 2005 from <http://www.securityfocus.com/infocus/1520>
- Karger, P. A., & Schell, R. R. (2002). "Thirty Years Later: Lessons from the MULTICS Security Evaluation," *Proceedings of the 2002 Annual Computer Security Applications Conference*, Las Vegas NV.
- , (1974). *MULTICS Security Evaluation: Vulnerability Analysis*. ESD-TR-74-193. Hanscom AFB, Bedford MA: ESD/AFSC.

- Leedy, P. D., & Ormrod, J. E. (2001). *Practical Research: Planning and Design* (7th ed.). New Jersey: Merrill Prentice Hall.
- Lundin, E., & Jonsson, E. (2002). *Survey of Intrusion Detection Research*. Technical Report No. 02-04. Goteborg, Sweden: Chalmers University of Technology.
- Lunt, T. F. (1988). "Automated Audit Trail Analysis and Intrusion Detection: A Survey," *Proceedings of the 11th National Computer Security Conference*, Baltimore, MD, 65-73.
- Maconachy, W. V., Schou, C. D., Ragsdale, D., & Welch, D. (2001). "A Model for Information Assurance: An Integrated Approach," *Proceedings of the 2001 IEEE Workshop on Information Assurance and Security*, United States Military Academy, West Point NY, 306-310.
- Mason, R. O., McKenney, J. L., & Copeland, D. G. (1997a). "Developing an Historical Tradition in MIS Research," *MIS Quarterly*, 21(3), 257-278.
- (1997b). "An Historical Method for MIS Research: Steps and Assumptions," *MIS Quarterly*, 21(3), 307-320.
- McHugh, J. (2001). "Intrusion and Intrusion Detection" *International Journal of Information Security*, 1(1), 14-35.
- McKnight, W. L. (2002). "What is Information Assurance?," *Crosstalk: The Journal of Defense Software Engineering*, January 12, 2005 from <http://www.stsc.hill.af.mil/crosstalk/2002/07/mcknight.html>
- Myers, P. (1980). *Subversion: The Neglected Aspect of Computer Security*. Master's thesis. Monterey CA: Naval Postgraduate School
- National Computer Security Center. (1985). *Department of Defense Trusted Computer System Evaluation Criteria (Orange Book)*. DoD 5200.28-STD.
- National Security Agency. (2002). *Information Assurance Technical Framework* (Release 3.1) from http://www.iaf.net/framework_docs/version-3_1/index.cfm

- Neumann, P. G., Robinson, L., Levitt, K. N., Boyer, R. S., & Saxena, A. R. (1975). *A Provably Secure Operating System*. No. M79-225). Menlo Park CA: Stanford Research Institute.
- O'Brien, J., Remenyi, D., & Keaney, A. (2004). "Historiography-A Neglected Research Method in Business and Management Studies," *Electronic Journal of Business Research Methods*, 2(2), 135-144.
- Rogers, L. R. (2002). *Home Computer Security*. Retrieved February 2, 2004 from <http://www.cert.org/homeusers/HomeComputerSecurity/#4>
- Snapp, S. R., Brentano, J., Dias, G. V., Goan, T. L., Heberlein, L. T., & Ho, C. et al. (1992). "DIDS (Distributed Intrusion Detection System) - Motivation, Architecture, and an Early Prototype," *Proceedings of the Fifteenth National Computer Security Conference*, Baltimore MD.
- Vaccaro, H. S., & Liepins, G. E. (1989). "Detection of Anomalous Computer Session Activity," *Proceedings of the 1989 IEEE Symposium on Research in Security and Privacy*, Oakland CA, 280-289.
- Ware, W. H. (1970). *Security Controls for Computer Systems (U): Report of Defense Science Board Task Force on Computer Security*. Santa Monica CA: The RAND Corporation.
- Whitmore, J., Bensoussan, A., Green, P., Hunt, D., Robziar, A., & Stern, J. (1973). *Design for MULTICS Security Enhancements*. ESD-TR-77-176. Hanscom AFB, Bedford MA: ESD/AFSC.

Vita

James L. M. Hart was born in Gretna, Louisiana in 1976 to Jerry and Judy Hart. James graduated from West Jefferson High School in Harvey, Louisiana in 1994. In 1996, he enlisted in the U.S. Air Force and married shortly thereafter. Upon graduating from Basic Military Training at Lackland Air Force Base, San Antonio, Texas, James attended technical training at Keesler Air Force, Mississippi. In 1997, James reported to his first duty station at Holloman Air Force Base, Alamogordo, New Mexico where he served as a Ground Radio Communications Journeyman in the 49th Communications Squadron. James attended undergraduate courses while off-duty and, in 2000, received his Bachelor of Science Degree in Management/Computer Information Systems from Park University where he graduated Magna Cum Laude. Also in 2000, James was competitively selected to attend Air Force Officer Training School at Maxwell Air Force Base, Montgomery, Alabama. Upon successful completion of the course, James was commissioned as a second lieutenant in the Air Force and relocated to Tyndall Air Force Base, Panama City, Florida where he served in a variety of positions culminating as Executive Officer of the 325th Operations Group. During his tenure there, James was again competitively selected to attend a graduate program in Information Resource Management at the Air Force Institute of Technology (AFIT), Wright-Patterson Air Force Base, Dayton, Ohio. Following completion of the graduate program at AFIT, James will be assigned to the Headquarters Staff, Air Intelligence Agency, Lackland Air Force Base, San Antonio, Texas. James and his wife have two children.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 074-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of the collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 21-03-2005		2. REPORT TYPE Master's Thesis		3. DATES COVERED (From – To) June 2004 – March 2005	
4. TITLE AND SUBTITLE An Historical Analysis of Factors Contributing to the Emergence of the Intrusion Detection Discipline and its Role in Information Assurance				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Hart, James L.M., Capt, USAF				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way, Building 641 WPAFB OH 45433-8865				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GIR/ENV/05M-06	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Information Assurance Laboratory Information Resources Management College National Defense University, Attn: Lt Col Clifton Poole Fort McNair, Washington, D.C.				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT In 2003, Gartner, Inc., predicted the inevitable demise of the intrusion detection (ID) market, a major player in the computer security technology industry. In light of this prediction, IT executives need to know if intrusion detection technologies serve a strategic purpose within the framework of information assurance (IA). This research investigated the historical background and circumstances that led to the birth of the intrusion detection field and explored the evolution of the discipline through current research in order to identify appropriate roles for IDS technology within an information assurance framework. The research identified factors contributing to the birth of ID including increased procurement and employment of resource-sharing computer systems in the DoD, a growing need to operate in an open computing environment while maintaining security and the unmanageable volume of audit data produced as a result of security requirements. The research also uncovered six trends that could be used to describe the evolution of the ID discipline encompassing passive to active response mechanisms, centralized to distributed management platforms, centralized to distributed/agent-based detection, single to multiple detection approaches within a system, host-based to network to hybrid analysis and software-based to hardware-based/in-line devices. Finally, the research outlined three roles suitable for IDS to fulfill within the IA framework including employing IDS as a stimulus to incident response mechanisms, as a forensic tool for gathering evidence of computer misuse and as a vulnerability assessment or policy enforcement facility.					
15. SUBJECT TERMS Intrusion Detection, Historiography, Historical Research, Information Assurance, Defense In Depth, History					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Captain David D. Bouvin, Associate Professor of IRM
U	U	U	UU	81	19b. TELEPHONE NUMBER (Include area code) (937) 255-6565, ext 4742 (david.bouvin@afit.edu)

